

TWO RECURSIVE GMRES-TYPE METHODS FOR SHIFTED LINEAR SYSTEMS WITH GENERAL PRECONDITIONING*

KIRK M. SOODHALTER[†]

Abstract. We present two minimum residual methods for solving sequences of shifted linear systems, the right-preconditioned shifted GMRES and shifted Recycled GMRES algorithms. These methods are compatible with general preconditioning of all systems, and when restricted to right preconditioning, require no extra applications of the operator or preconditioner. These methods perform a minimum residual iteration for the base system while improving the approximations for the shifted systems at little additional cost. The iteration continues until the base system approximation is of satisfactory quality. The method is then recursively called for the remaining unconverged systems. We present both methods inside of a general framework which allows these techniques to be extended to the setting of flexible preconditioning and inexact Krylov methods. We present some analysis of such methods and numerical experiments demonstrating the effectiveness of the algorithms we have derived.

Key words. Krylov subspace methods, shifted linear systems, parameterized linear systems, quantum chromodynamics

AMS subject classifications. 65F10, 65F50, 65F08

1. Introduction. We develop techniques for solving a family (or a sequence of families) of linear systems in which the coefficient matrices differ only by a scalar multiple of the identity. There are many applications which warrant the solution of a family of shifted linear systems, such as those arising in lattice quantum chromodynamics (QCD) (see, e.g., [9]) as well as other applications such as Tikhonov-Philips regularization, global methods of nonlinear analysis, and Newton trust region methods [3]. The goal is to develop a framework in which minimum residual methods can be applied to shifted systems in a way that:

- (a) allows us to exploit the relationships between the coefficient matrices
- (b) is compatible with general (right) preconditioning.

Consider a family of shifted linear systems, which we parameterize by ℓ , i.e.,

$$\left(\mathbf{A} + \sigma^{(\ell)}\mathbf{I}\right)\mathbf{x}^{(\ell)} = \mathbf{b} \quad \text{for } \ell = 1, \dots, L. \quad (1.1)$$

We call the numbers $\{\sigma^{(\ell)}\}_{\ell=1}^L \subset \mathbb{C}$ *shifts*, \mathbf{A} the *base matrix*, and $\mathbf{A} + \sigma\mathbf{I}$ a *shifted matrix*. Systems of the form (1.1) are called *shifted linear systems*. Krylov subspace methods have been proposed to simultaneously solve this family of systems, see, e.g., [7, 8, 26]. These methods satisfy requirement (a) but are not compatible with general preconditioning strategies, as they rely on the invariance of the Krylov subspace under constant shift of the coefficient matrix; cf. (2.5).

We can introduce an additional parameter i , which indexes a sequence of matrices $\{\mathbf{A}_i\} \subset \mathbb{C}^{n \times n}$, and for each i , we solve a family of the form

$$\left(\mathbf{A}_i + \sigma_i^{(\ell)}\mathbf{I}\right)\mathbf{x}_i^{(\ell)} = \mathbf{b}_i \quad \text{for } \ell = 1 \dots L_i \quad (1.2)$$

*This version dated March 19, 2014.

[†]Industrial Mathematics Institute, Johannes Kepler University, Altenbergerstraße 69, A-4040 Linz, Austria. (kirk.soodhalter@indmath.uni-linz.ac.at)

Augmented Krylov subspace methods have been proposed for efficiently solving a sequence of linear systems with a slowly changing coefficient matrix, allowing important spectral information generated while solving $\mathbf{A}_i \mathbf{x}_i = \mathbf{b}_i$ to be used to augment the Krylov subspace generated when solving $\mathbf{A}_{i+1} \mathbf{x}_{i+1} = \mathbf{b}_{i+1}$; see, e.g., [19, 24, 31]. In cases such as a Newton iteration, these matrices are available one at a time, while in a case such as an implicit time-stepping scheme, the matrix may not change at all.

In [28], the authors explored solving a family of shifted systems over an augmented Krylov subspace. Specifically, the goal was to develop a method which solved the family of systems simultaneously, using one augmented subspace to extract all candidate solutions, which also had a fixed storage requirement, independent of the number of shifts L . It was shown that in general within the framework of GMRES for shifted systems [8] and subspace recycling [19], such a method, *does not exist*.

A conclusion one can draw from [28] is that we should consider avoiding methods relying on the invariance of Krylov subspaces under a constant shift of the identity; cf. (2.5). Relying on this invariance imposes restrictions on our ability to develop an algorithm. Furthermore, relying on this shift invariance means we cannot use arbitrary preconditioners, an unfortunate limitation.

Learning from the results in [28], we focus on methods which do not rely on the shift invariance. Rather than focusing on specific Krylov subspace techniques (augmented or not), we instead begin by developing a general framework of minimum residual projection techniques for shifted linear systems. In this framework, we extract candidate solutions for all shifted systems from one subspace and we select each candidate solution according to a minimum residual Petrov-Galerkin condition. This framework is compatible with arbitrary right preconditioners, and the computational cost for each additional shifted system is small. By specifying subspaces once the framework is developed, we derive GMRES and recycled GMRES (rGMRES) methods for shifted systems that are compatible with general right preconditioning. Though not considered in this paper, the framework is also compatible with flexible and inexact Krylov methods.

In this work, we restrict ourselves to right preconditioned methods. In theory, everything we develop is compatible with left or two-sided preconditioning; but by restricting to right preconditioning, we derive methods which require extra storage but no extra applications of the operator or preconditioner. Furthermore, in a right preconditioned minimum residual method, we are minimizing the unpreconditioned residual 2-norm rather than in some other norm; see [25], for more details.

The rest of this paper is organized as follows. In the next section, we review the minimum residual Krylov subspace method GMRES as well as two GMRES variants, one for shifted linear systems and the other extending GMRES to the augmented Krylov subspace setting, i.e., Recycled GMRES. In Section 3, we present a general framework to perform minimum residual projections of the shifted system residuals with respect to the search space generated for the base system. In Subsection 3.1 we use this framework to derive our **shifted GMRES** method and in Subsection 3.2 we derive a **shifted Recycled GMRES** method. In Section 4, we present some analysis of the expected performance of methods derived from this framework. In Section 5, we present some numerical results.

2. Preliminaries. We begin with a brief review of Krylov subspace methods as well as techniques of subspace recycling and for solving shifted linear system. Recall that in many Krylov subspace iterative methods for solving

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{2.1}$$

with $\mathbf{A} \in \mathbb{C}^{n \times n}$, we generate an orthonormal basis for

$$\mathcal{K}_j(\mathbf{A}, \mathbf{u}) = \text{span} \{ \mathbf{u}, \mathbf{A}\mathbf{u}, \dots, \mathbf{A}^{j-1}\mathbf{u} \} \quad (2.2)$$

with the Arnoldi process, where \mathbf{u} is some starting vector. Let $\mathbf{V}_j \in \mathbb{C}^{n \times j}$ be the matrix with orthonormal columns generated by the Arnoldi process spanning $\mathcal{K}_j(\mathbf{A}, \mathbf{u})$. Then we have the Arnoldi relation

$$\mathbf{A}\mathbf{V}_j = \mathbf{V}_{j+1}\bar{\mathbf{H}}_j \quad (2.3)$$

with $\bar{\mathbf{H}}_j \in \mathbb{C}^{(j+1) \times j}$; see, e.g., [22, Section 6.3] and [27]. Let \mathbf{x}_0 be an initial approximation to the solution of a linear system we wish to solve and $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$ be the initial residual. At iteration j , we choose $\mathbf{x}_j = \mathbf{x}_0 + \mathbf{t}_j$, with $\mathbf{t}_j \in \mathcal{K}_j(\mathbf{A}, \mathbf{r}_0)$. In GMRES [23], \mathbf{t}_j satisfies

$$\mathbf{t}_j = \underset{\mathbf{t} \in \mathcal{K}_j(\mathbf{A}, \mathbf{r}_0)}{\text{argmin}} \|\mathbf{b} - \mathbf{A}(\mathbf{x}_0 + \mathbf{t})\|,$$

which is equivalent to solving the smaller minimization problem

$$\mathbf{y}_j = \underset{\mathbf{y} \in \mathbb{C}^j}{\text{argmin}} \left\| \bar{\mathbf{H}}_j \mathbf{y} - \|\mathbf{r}_0\| \mathbf{e}_1^{(j+1)} \right\|, \quad (2.4)$$

where $\mathbf{e}_1^{(i)}$ denotes the i th Cartesian basis vector in \mathbb{C}^i . We then set $\mathbf{x}_j = \mathbf{x}_0 + \mathbf{V}_j \mathbf{y}_j$. Recall that in restarted GMRES, often called GMRES(m), we run an m -step cycle of the GMRES method and compute an approximation \mathbf{x}_m . We halt the process, discard \mathbf{V}_m , and restart with the new residual. This process is repeated until we achieve convergence. An adaption of restarted GMRES to solve (1.1) has been previously proposed; see, e.g., [8].

Many methods for the simultaneous solution of shifted systems (see, e.g., [5, 7, 8, 9, 14, 26]) take advantage of the fact that for any shift $\sigma \in \mathbb{C}$, the Krylov subspace generated by \mathbf{A} and \mathbf{b} is invariant under the shift, i.e.,

$$\mathcal{K}_j(\mathbf{A}, \mathbf{b}) = \mathcal{K}_j(\mathbf{A} + \sigma \mathbf{I}, \tilde{\mathbf{b}}), \quad (2.5)$$

as long as the starting vectors are collinear, i.e., $\tilde{\mathbf{b}} = \beta \mathbf{b}$ for some $\beta \in \mathbb{C}$, with a shifted Arnoldi relation similar to (2.3)

$$(\mathbf{A} + \sigma \mathbf{I})\mathbf{V}_j = \mathbf{V}_{j+1}\bar{\mathbf{H}}_j(\sigma), \quad (2.6)$$

where $\bar{\mathbf{H}}_j(\sigma) = \bar{\mathbf{H}}_j + \sigma \begin{bmatrix} \mathbf{I}_{m \times m} \\ \mathbf{0}_{1 \times m} \end{bmatrix}$. This collinearity must be maintained at restart. In [28], this was shown to be a troublesome restriction when attempting to extend such techniques augmented Krylov methods. Furthermore, note that the shift-invariance no longer holds if general preconditioning is used; e.g., in the case of right preconditioning

$$\mathcal{K}_j(\mathbf{A}\mathbf{M}^{-1}, \mathbf{b}) = \mathcal{K}_j((\mathbf{A} + \sigma \mathbf{I})\mathbf{M}^{-1}, \tilde{\mathbf{b}})$$

does **not** hold in general. It should be pointed out that specific polynomial preconditioners can be constructed (see, e.g., [1, 13, 32]) for which shift invariance can be maintained. The invariance (2.5) can lead to great savings in memory costs, since one needs only generate a basis for a single Krylov subspace to solve multiple linear

systems. However, because of the restrictions imposed, we avoid basing our method on it, while still exploiting the relationships between the shifted matrices.

We briefly review the Recycled GMRES method (sometimes referred to as rGMRES for brevity) described in [19]. Deflation and augmentation techniques designed specifically for symmetric linear systems have also been proposed; see, e.g., [24, 31]. For a more general framework for these types of methods, see [10]. Gaul and Schlömer describe recycling techniques in the context of self-adjoint operator equations in a general Hilbert space [11].

The Recycled GMRES algorithm represents the confluence of two approaches: those descending from the implicitly restarted Arnoldi method [15], such as Morgan's GMRES-DR [16], and those descending from de Sturler's GCRO method [29]. GMRES-DR is a restarted GMRES algorithm, where at the end of each cycle, harmonic Ritz vectors are computed, and a subset of them are used to augment the Krylov subspace generated at the next cycle. The GCRO method allows the user to select the optimal correction over arbitrary subspaces. This concept is extended by de Sturler in [30], where a framework is provided for selecting the optimal subspace to retain from one cycle to the next so as to minimize the error produced by discarding useful information accumulated in the subspace for candidate solutions before restart. This algorithm is called GCROT, where OT stands for optimal truncation. A simplified version of the GCROT approach, based on restarted GMRES (called LGMRES) is presented in [2]. Parks et al. in [19] combine the ideas of [16] and [30] and extend them to a sequence of slowly-changing linear systems. They call their method GCRO-DR (which we will denote Recycled GMRES).

Suppose we are solving (2.1), and we have a k -dimensional subspace \mathcal{U} whose image under the action of \mathbf{A} is $\mathcal{C} = \mathbf{A}\mathcal{U}$. Let \mathbf{P} be the orthogonal projector onto \mathcal{C}^\perp . Let \mathbf{x}_0 be such that $\mathbf{r}_0 \in \mathcal{C}^\perp$ (this is always cheaply available). We generate the Krylov subspace with respect to the projected operator \mathbf{PA} , $\mathcal{K}_m(\mathbf{PA}, \mathbf{r}_0)$. At iteration m , the Recycled GMRES method generates the approximation

$$\mathbf{x}_m = \mathbf{x}_0 + \mathbf{s}_m + \mathbf{t}_m$$

where $\mathbf{s}_m \in \mathcal{U}$ and $\mathbf{t}_m \in \mathcal{K}_m(\mathbf{PA}, \mathbf{r}_0)$. The corrections \mathbf{s}_m and \mathbf{t}_m are chosen according to the minimum residual, Petrov-Galerkin condition over the augmented Krylov subspace, i.e.,

$$\mathbf{r}_m \perp \mathbf{A}(\mathcal{U} + \mathcal{K}_m(\mathbf{PA}, \mathbf{r}_0)). \quad (2.7)$$

At the end of the cycle, an updated \mathcal{U} is constructed, the Krylov subspace basis is discarded, and we restart. At convergence, \mathcal{U} is saved, to be used when solving the next linear system.

In terms of implementation, Recycled GMRES can be described as a modification of the GMRES algorithm. Let $\mathbf{U} \in \mathbb{C}^{n \times k}$ have columns spanning \mathcal{U} , scaled such that $\mathbf{C} = \mathbf{AU}$ has orthonormal columns. Then we can explicitly construct $\mathbf{P} = \mathbf{I} - \mathbf{CC}^*$. At each iteration, applying \mathbf{P} is equivalent to performing k steps of the Modified Gram-Schmidt process to orthogonalize the new Arnoldi vector against the columns of \mathbf{C} . The orthogonalization coefficients generated at step m are stored in the m th column of $\mathbf{B}_m = \mathbf{C}^* \mathbf{A} \mathbf{V}_m$, and \mathbf{B}_{m+1} is simply \mathbf{B}_m with one new column appended. Let $\bar{\mathbf{H}}_m$ and \mathbf{V}_m be defined as before, but for the projected Krylov subspace $\mathcal{K}_m(\mathbf{PA}, \mathbf{r}_0)$. Enforcing (2.7) is equivalent to solving the GMRES minimization problem (2.4) for $\mathcal{K}_m(\mathbf{PA}, \mathbf{r}_0)$ and setting

$$\mathbf{s}_m = -\mathbf{UB}_m \mathbf{y}_m \quad \text{and} \quad \mathbf{t}_m = \mathbf{V}_m \mathbf{y}_m,$$

so that

$$\mathbf{x}_m = \mathbf{x}_0 - \mathbf{U}\mathbf{B}_m\mathbf{y}_m + \mathbf{V}_m\mathbf{y}_m = \mathbf{x}_0 + [\mathbf{U} \quad \mathbf{V}_m] \begin{bmatrix} -\mathbf{B}_m\mathbf{y}_m \\ \mathbf{y}_m \end{bmatrix}.$$

This is a consequence of the fact that the Recycled GMRES least squares problem, as stated in [19, Equation 2.13] can be satisfied exactly in the first k rows.

3. A direct projection framework. We develop a general framework of minimum residual methods for shifted linear systems which encompasses both unpreconditioned and preconditioned systems. For simplicity, we restrict our description to two model problems: the unpreconditioned problem

$$\mathbf{A}\mathbf{x} = \mathbf{b} \quad \text{and} \quad (\mathbf{A} + \sigma\mathbf{I})\mathbf{x}(\sigma) = \mathbf{b} \quad (3.1)$$

and the right-preconditioned problem

$$\mathbf{A}\mathbf{M}^{-1}\mathbf{w} = \mathbf{b} \quad \text{and} \quad (\mathbf{A} + \sigma\mathbf{I})\mathbf{M}^{-1}\mathbf{w}(\sigma) = \mathbf{b} \quad (3.2)$$

where $\mathbf{w}_0 = \mathbf{M}\mathbf{x}_0$ and $\mathbf{w}_0(\sigma) = \mathbf{M}\mathbf{x}_0(\sigma)$, and after m iterations we set $\mathbf{x}_m = \mathbf{M}^{-1}\mathbf{w}_m$ and we set $\mathbf{x}_m(\sigma) = \mathbf{M}^{-1}\mathbf{y}_m(\sigma)$. However, what we develop is fully applicable to sequences of linear systems of the form (1.2).

We describe these techniques in terms of a general sequence of nested subspaces

$$\mathbf{S}_1 \subset \mathbf{S}_2 \subset \cdots \mathbf{S}_m \subset \cdots$$

This allows us to cleanly present these techniques as minimum residual projection methods and later to give clear analysis, applicable to any method fitting into this framework. Then we can derive different methods by specifying \mathbf{S}_m , e.g., $\mathbf{S}_m = \mathcal{K}_m(\mathbf{A}, \mathbf{r}_0)$.

Let $\{\mathbf{S}_m\}_{i=1}^m$ be the nested sequence of subspaces produced by some iterative method for solving (3.1) or (3.2), after m iterations. In the unpreconditioned case (3.1), suppose we have initial approximations \mathbf{x}_0 and $\mathbf{x}_0(\sigma)$ for the base and shifted systems, respectively. For conciseness, let us denote $\mathbf{A}(\sigma) = \mathbf{A} + \sigma\mathbf{I}$. At iteration m , we compute corrections $\mathbf{t}_m, \mathbf{t}_m(\sigma) \in \mathbf{S}_m$ which satisfy the minimum residual conditions

$$\mathbf{b} - \mathbf{A}(\mathbf{x}_0 + \mathbf{t}_m) \perp \mathbf{A}\mathbf{S}_m \quad \text{and} \quad \mathbf{b} - \mathbf{A}(\sigma)(\mathbf{x}_0(\sigma) + \mathbf{t}_m(\sigma)) \perp \mathbf{A}(\sigma)\mathbf{S}_m. \quad (3.3)$$

In the preconditioned case (3.2), suppose we begin with initial approximations $\mathbf{w}_0 = \mathbf{M}\mathbf{x}_0$ and $\mathbf{w}_0(\sigma) = \mathbf{M}\mathbf{x}_0(\sigma)$. Let us denote the preconditioned operators

$$\mathbf{A}_p = \mathbf{A}\mathbf{M}^{-1} \quad \text{and} \quad \mathbf{A}_p(\sigma) = (\mathbf{A} + \sigma\mathbf{I})\mathbf{M}^{-1}$$

At iteration m , we compute corrections $\mathbf{t}_m, \mathbf{t}_m(\sigma) \in \mathbf{S}_m$ which satisfy the minimum residual conditions

$$\mathbf{b} - \mathbf{A}_p(\mathbf{w}_0 + \mathbf{t}_m) \perp \mathbf{A}_p\mathbf{S}_m \quad \text{and} \quad \mathbf{b} - \mathbf{A}_p(\sigma)(\mathbf{w}_0(\sigma) + \mathbf{t}_m(\sigma)) \perp \mathbf{A}_p(\sigma)\mathbf{S}_m. \quad (3.4)$$

We emphasize that *the same preconditioner* is used for all systems.

In this framework, we assume that the Galerkin condition for the base case is applied via a predefined iterative method, the method which generates the sequence $\{\mathbf{S}_m\}$. Therefore, it suffices to describe the residual projection for the shifted system. We can write the update of the shifted system approximation by explicitly constructing the orthogonal projector which is applied during a Petrov-Galerkin projection.

Let $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_m\}$ be a basis for \mathcal{S}_m which we take as the columns of $\mathbf{S}_m \in \mathbb{C}^{n \times m}$. Then we can write this projection and update

$$\begin{aligned} \mathbf{r}_m(\sigma) &= \mathbf{r}_0(\sigma) - \mathbf{A}(\sigma)\mathbf{S}_m\mathbf{y}_m(\sigma) \quad \text{and} \\ \mathbf{x}_m(\sigma) &= \mathbf{x}_0(\sigma) + \mathbf{S}_m\mathbf{y}_m(\sigma) \end{aligned} \quad (3.5)$$

where $\mathbf{y}_m(\sigma) = \mathbf{N}_m(\sigma)^{-1} (\mathbf{A}(\sigma)\mathbf{S}_m)^* \mathbf{r}_0(\sigma)$ and $\mathbf{N}_m(\sigma) = (\mathbf{S}_m^* \mathbf{A}(\sigma)^* \mathbf{A}(\sigma) \mathbf{S}_m)$ is the projection scaling matrix, since we assume that $\mathbf{A}(\sigma)\mathbf{S}_m$ does not have orthonormal columns. For an arbitrary sequence of subspaces, constructing such projectors could be quite costly. However, for special choices (i.e., Krylov subspaces), the projections can be built from already-computed quantities.

In the following subsections, we derive new methods by specifying subspaces $\{\mathcal{S}_m\}$ and a matrix \mathbf{S}_m which has as columns a basis for \mathcal{S}_m . This will define our projection scaling matrix $\mathbf{N}_m(\sigma)$. For each \mathcal{S}_m we then show that $\mathbf{N}_m(\sigma)$ is composed of blocks which can be built from already-computed quantities. Thus, for appropriate choices of \mathcal{S}_m , the projection, i.e., either (3.3) or (3.4), can be inexpensively applied.

We highlight that a strength of this framework is the ease with which it can be used to develop a method for shifted systems on top of an existing iterative method, with few modifications. As the framework only requires a sequence of nested linear systems, it is completely compatible with both standard Krylov subspace methods as well as flexible and inexact Krylov subspace methods.

3.1. A GMRES method for shifted systems. In the case that we apply the GMRES iteration to the base system, at iteration m , our search space $\mathcal{S}_m \leftarrow \mathcal{K}_m(\mathbf{A}, \mathbf{r}_0)$ is the m th Krylov subspace. The matrix $\mathbf{S}_m \leftarrow \mathbf{V}_m$ has the first m Arnoldi vectors as columns. The projection and update (3.5) can be simplified due to the shifted Arnoldi relation (2.6). The matrix $\mathbf{N}_m(\sigma) \leftarrow \overline{\mathbf{H}}_m(\sigma)^* \overline{\mathbf{H}}_m(\sigma) \in \mathbb{C}^{m \times m}$ can be constructed from the already computed upper Hessenberg matrix. Thus the projection (3.3) can be rewritten

$$\begin{aligned} \mathbf{r}_m(\sigma) &= \mathbf{r}_0(\sigma) - \mathbf{V}_{m+1} \overline{\mathbf{H}}_m(\sigma) \mathbf{y}_m(\sigma) \quad \text{and} \\ \mathbf{x}_m(\sigma) &= \mathbf{x}_0(\sigma) - \mathbf{V}_m(\sigma) \mathbf{y}_m(\sigma) \end{aligned}$$

where $\mathbf{y}_m(\sigma) \leftarrow (\overline{\mathbf{H}}(\sigma)^* \overline{\mathbf{H}}(\sigma))^{-1} \overline{\mathbf{H}}_m(\sigma)^* \mathbf{V}_{m+1}^* \mathbf{r}_0(\sigma)$. As it can be appreciated, applying this Petrov-Gallerkin condition is equivalent to solving the least squares problem

$$\mathbf{y}_m(\sigma) = \underset{\mathbf{y} \in \mathbb{C}^i}{\operatorname{argmin}} \left\| \overline{\mathbf{H}}_m(\sigma) - \mathbf{V}_{m+1}^* \mathbf{r}_0(\sigma) \right\|$$

and setting $\mathbf{x}_m(\sigma) = \mathbf{x}_0(\sigma) - \mathbf{V}_m \mathbf{y}_m(\sigma)$. It should be noted that this method differs from the GMRES method for shifted systems of Frommer and Glässner [8], which is derived from the invariance (2.5).

3.1.1. Preconditioning. Introducing preconditioning into this setting presents complications. No longer can we use the shifted Arnoldi relation (2.6) as we could in the unpreconditioned case. However, by storing some extra vectors, we can apply the same Petrov-Gallerkin condition with no additional application of the operator or preconditioner.

Recall that in right-preconditioned GMRES (see, e.g., [22, Sections 9.3.2 and 9.4.1]) that the $\mathcal{S}_m \leftarrow \mathbf{M}^{-1} \mathcal{K}(\mathbf{A}_p, \mathbf{r}_0)$, and $\mathbf{S}_m \leftarrow \mathbf{M}^{-1} \mathbf{V}_m$. This space is never explicitly constructed, though, since if \mathbf{y}_m is the solution to the GMRES least squares problem (2.4) in the preconditioned case, we simply set $\mathbf{x}_m = \mathbf{x}_0 + \mathbf{M}^{-1} (\mathbf{V}_m \mathbf{y}_m)$.

However, in the right-preconditioned Arnoldi process, one can store a basis for the subspace $\mathbf{M}^{-1}\mathcal{K}(\mathbf{A}_p, \mathbf{r}_0)$ at no additional computational cost. For all $1 \leq i \leq m$, let $\mathbf{z}_m = \mathbf{M}^{-1}\mathbf{v}_m$, and let these vectors be the columns of $\mathbf{Z}_m \in \mathbb{C}^{n \times i}$ so that $\mathbf{Z}_m = \mathbf{M}^{-1}\mathbf{V}_m$.

With these vectors, we can perform the minimum residual projection for the shifted system such that it requires no additional applications of the operator or the preconditioner. Observe that we can write $\mathbf{A}_p(\sigma) = \mathbf{A}_p + \sigma\mathbf{M}^{-1}$. As in the unpreconditioned case, we explicitly project the residual, but this time orthogonal to $\mathbf{A}_p(\sigma)\mathcal{K}(\mathbf{A}_p, \mathbf{r}_0)$,

$$\mathbf{r}_m(\sigma) = \mathbf{r}_0(\sigma) - (\mathbf{A}_p + \sigma\mathbf{M}^{-1}) \mathbf{V}_m \mathbf{N}_m(\sigma)^{-1} [(\mathbf{A}_p + \sigma\mathbf{M}^{-1}) \mathbf{V}_m]^* \mathbf{r}_0(\sigma) \quad (3.6)$$

where $\mathbf{N}_m(\sigma) \leftarrow [(\mathbf{A}_p + \sigma\mathbf{M}^{-1}) \mathbf{V}_m]^* [(\mathbf{A}_p + \sigma\mathbf{M}^{-1}) \mathbf{V}_m]$. We can then use the right-preconditioned shifted Arnoldi relation

$$(\mathbf{A}_p + \sigma\mathbf{M}^{-1}) \mathbf{V}_m = \mathbf{V}_{m+1} \bar{\mathbf{H}}_m + \sigma \mathbf{Z}_m$$

to rewrite

$$\mathbf{N}_m(\sigma) = \bar{\mathbf{H}}_m^* \bar{\mathbf{H}}_m + \sigma \bar{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{Z}_m + \sigma \mathbf{Z}_m^* \mathbf{V}_{m+1} \bar{\mathbf{H}}_m + |\sigma|^2 \mathbf{Z}_m^* \mathbf{Z}_m.$$

Thus, the approximation update and the residual projection (3.6) can be rewritten

$$\begin{aligned} \mathbf{x}_m(\sigma) &= \mathbf{x}_0(\sigma) + \mathbf{Z}_m \mathbf{y}_m(\sigma) \\ \mathbf{r}_m(\sigma) &= \mathbf{r}_0(\sigma) - (\mathbf{V}_{m+1} \bar{\mathbf{H}}_m + \sigma \mathbf{Z}_m) \mathbf{y}_m(\sigma). \end{aligned}$$

where $\mathbf{y}_m(\sigma) = \mathbf{N}_m(\sigma)^{-1} [(\mathbf{V}_{m+1} \bar{\mathbf{H}}_m + \sigma \mathbf{Z}_m)]^* \mathbf{r}_0(\sigma)$. This projection process involves only matrices ($\bar{\mathbf{H}}_m$, \mathbf{V}_{m+1} , and \mathbf{Z}_{m+1}) which have already been computed in the course of right-preconditioned GMRES for the base system. The matrices $\bar{\mathbf{H}}_m^* \bar{\mathbf{H}}_m$, $\bar{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{Z}_m$, and $\mathbf{Z}_m^* \mathbf{Z}_m$ can be computed once, regardless of the number of shifted systems. The only operation whose cost depends on the number of shifts is the solving of a linear system for $\mathbf{N}_m(\sigma)$, which must be performed for each σ . This solution of a symmetric $m \times m$ linear system costs $\mathcal{O}(m^3)$ flops. The right-preconditioned shifted GMRES algorithm (sometimes referred to as sGMRES) is shown in Algorithm 3.1. Observe that an implementation can rely heavily on an existing GMRES code. It should be noted that all but one step of the shifted residual projections can be formulated in terms of block/BLAS-3 operations so that computations for all shifts are performed simultaneously. Only the computation of $\mathbf{y}_m(\sigma)$ must be performed for each shift separately, in a loop, as it involves $\mathbf{N}_m(\sigma)$.

3.2. An rGMRES method for shifted systems. Suppose now that our iteration for the base system is Recycled GMRES.

We begin by projecting the initial residual $\mathbf{r}_{-1}(\sigma)$ associated to initial approximation $\mathbf{x}_{-1}(\sigma)$, so that we begin with $\mathbf{r}_0(\sigma) \perp \mathbf{A}(\sigma)\mathcal{U}$. This is equivalent to computing the minimum residual correction $\mathbf{t}_0(\sigma) \in \mathcal{U}$ and setting $\mathbf{x}_0(\sigma) = \mathbf{x}_{-1}(\sigma) + \mathbf{t}_0(\sigma)$. In Recycled GMRES, such a projection is necessary to correctly derive the algorithm. For the shifted system, the projection is not necessary, but it does allow for some simplifications later in the derivation,

$$\mathbf{x}_0(\sigma) = \mathbf{x}_{-1}(\sigma) + \mathbf{U} \mathbf{y}_0(\sigma) \quad \text{and} \quad \mathbf{r}_0(\sigma) = \mathbf{r}_{-1}(\sigma) - \mathbf{A}(\sigma) \mathbf{U} \mathbf{y}_0(\sigma), \quad (3.7)$$

where $\mathbf{y}_0(\sigma) = \mathbf{N}_0(\sigma)^{-1}(\mathbf{A}(\sigma)\mathbf{U})^* \mathbf{r}_{-1}$ and $\mathbf{N}_0(\sigma) = (\mathbf{A}(\sigma)\mathbf{U})^* (\mathbf{A}(\sigma)\mathbf{U})$. Since $\mathbf{A}(\sigma)\mathbf{U} = \mathbf{C} + \sigma\mathbf{U}$, this projection can be simplified and cheaply computed,

$$\mathbf{r}_0 = \mathbf{r}_{-1} - (\mathbf{C} + \sigma\mathbf{U})\mathbf{N}_m(\sigma)^{-1}(\mathbf{C} + \sigma\mathbf{U})^* \mathbf{r}_{-1}$$

where we rewrite $\mathbf{N}_0(\sigma) = \mathbf{I}_{k \times k} + \sigma\mathbf{C}^*\mathbf{U} + \bar{\sigma}\mathbf{U}^*\mathbf{C} + |\sigma|^2 \mathbf{U}^*\mathbf{U}$. The matrices $\mathbf{C}^*\mathbf{U}$ and $\mathbf{U}^*\mathbf{U}$ must only be computed once, regardless of the number of shifts, and for each shift, the solution of $\mathbf{N}_0(\sigma)\mathbf{y}_0(\sigma) = (\mathbf{C} + \sigma\mathbf{U})^* \mathbf{r}_{-1}(\sigma)$ must be computed.

After a cycle of Recycled GMRES for the base system, the projection and update (3.3) must be applied to each shifted system. At iteration m , our search space $\mathcal{S}_m \leftarrow \mathcal{U} + \mathcal{K}(\mathbf{P}\mathbf{A}_p, \mathbf{r}_0)$ is the m th augmented Krylov subspace. The augmented matrix $\mathbf{S}_m \leftarrow [\mathbf{U} \quad \mathbf{V}_m]$ contains as columns the basis for \mathcal{U} and $\mathcal{K}_m(\mathbf{P}\mathbf{A}_p, \mathbf{r}_0)$. In this case, we have $\mathbf{N}_m(\sigma) \leftarrow \{(\mathbf{A} + \sigma\mathbf{I}) [\mathbf{U} \quad \mathbf{V}_m]\}^* \{(\mathbf{A} + \sigma\mathbf{I}) [\mathbf{U} \quad \mathbf{V}_m]\}$. From [28], we have the identity

$$(\mathbf{A} + \sigma\mathbf{I}) [\mathbf{U} \quad \mathbf{V}_m] = [(\mathbf{C} + \sigma\mathbf{U}) \quad (\mathbf{C}\mathbf{B}_m + \mathbf{V}_{m+1}\bar{\mathbf{H}}_m + \sigma\mathbf{V}_m)].$$

Thus, in the unpreconditioned case, for the augmented Krylov subspace, we can rewrite (3.3)

$$\begin{aligned} \mathbf{r}_m(\sigma) &= \mathbf{r}_0(\sigma) - [(\mathbf{C} + \sigma\mathbf{U}) \quad (\mathbf{C}\mathbf{B}_m + \mathbf{V}_{m+1}\bar{\mathbf{H}}_m + \sigma\mathbf{V}_m)] \mathbf{y}_m(\sigma) \quad \text{and} \\ \mathbf{x}_m(\sigma) &= \mathbf{x}_0(\sigma) + [\mathbf{U} \quad \mathbf{V}_m] \mathbf{y}_m(\sigma) \end{aligned} \quad (3.8)$$

Algorithm 3.1: Right preconditioned shifted GMRES (**sGMRES()**)

Input : $\mathbf{A} \in \mathbb{C}^{n \times n}$; $\mathbf{b} \in \mathbb{C}^n$; $\{\sigma_\ell\}_{\ell=1}^L \subset \mathbb{C}$; Initial Approximations $\{\mathbf{x}(\sigma_\ell)\}_{\ell=1}^L$;
 $\varepsilon > 0$; Cycle length $m \in \mathbb{N}$
Output: $\{\mathbf{x}(\sigma_\ell)\}_{\ell=1}^L$ such that $\|\mathbf{r}(\sigma_\ell)\| / \|\mathbf{r}_0(\sigma_\ell)\| \leq \varepsilon$ for all ℓ

```

1 for  $\ell = 1 \dots L$  do
2    $\mathbf{r}(\sigma_\ell) = \mathbf{b} - (\mathbf{A} + \sigma_\ell \mathbf{I})\mathbf{x}(\sigma_\ell)$ 
3    $\gamma_1 = \|\mathbf{r}(\sigma_1)\|$ 
4   if  $L > 1$  then
5     while  $\|\mathbf{r}(\sigma_1)\| / \gamma_1 > \varepsilon$  do
6       Compute and overwrite  $\mathbf{x}(\sigma_1)$ ,  $\mathbf{r}(\sigma_1)$ ,  $\mathbf{V}_{m+1}$ ,  $\mathbf{Z}_m$ ,  $\bar{\mathbf{H}}_m$  by calling
       GMRES() for  $\mathbf{A} + \sigma_1 \mathbf{I}$ ,  $\mathbf{M}$ ,  $\mathbf{b}$ ,  $\mathbf{x}(\sigma_1)$ , and  $m$ 
7       Compute and overwrite  $\bar{\mathbf{H}}_m^* \bar{\mathbf{H}}_m$ ,  $\bar{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{Z}_m$ , and  $\mathbf{Z}_m^* \mathbf{Z}_m$ 
8       for  $\ell = 2 \dots s$  do
9          $\mathbf{N} \leftarrow \bar{\mathbf{H}}_m^* \bar{\mathbf{H}}_m + \sigma \bar{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{Z}_m + \bar{\sigma} \mathbf{Z}_m^* \mathbf{V}_{m+1} \bar{\mathbf{H}}_m + |\sigma|^2 \mathbf{Z}_m^* \mathbf{Z}_m$ 
10         $\mathbf{y} \leftarrow \mathbf{N}^{-1} [(\mathbf{V}_{m+1} \bar{\mathbf{H}}_m + \sigma \mathbf{Z}_m)]^* \mathbf{r}_0(\sigma)$ 
11         $\mathbf{x}(\sigma_\ell) \leftarrow \mathbf{x}_0(\sigma_\ell) + \mathbf{Z}_m \mathbf{y}$ 
12         $\mathbf{r}(\sigma_\ell) \leftarrow \mathbf{r}_0(\sigma_\ell) - (\mathbf{V}_{m+1} \bar{\mathbf{H}}_m + \sigma \mathbf{Z}_m) \mathbf{y}$ 
13      For all  $\ell = 2, \dots, L$  compute and overwrite  $\mathbf{x}(\sigma_\ell)$  by recursively calling
      sGMRES() for  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $\mathbf{M}$ ,  $\{\sigma_\ell\}_{\ell=2}^L$ ,  $\{\mathbf{x}(\sigma_\ell)\}_{\ell=2}^L$ ,  $\varepsilon$ , and  $m$ 
14   else
15     while  $\|\mathbf{r}(\sigma_1)\| / \gamma_1 > \varepsilon$  do
16       Compute and overwrite  $\mathbf{x}(\sigma_1)$ ,  $\mathbf{r}(\sigma_1)$  by calling GMRES() for  $\mathbf{A} + \sigma_1 \mathbf{I}$ ,  $\mathbf{M}$ ,
        $\mathbf{b}$ ,  $\mathbf{x}(\sigma_1)$ , and  $m$ 
```

where $\mathbf{y}_m(\sigma) = \mathbf{N}_m(\sigma)^{-1} \left[(\mathbf{C} + \sigma \mathbf{U}) \quad (\mathbf{C} \mathbf{B}_m + \mathbf{V}_{m+1} \bar{\mathbf{H}}_m^{(\sigma)}) \right]^* \mathbf{r}_0(\sigma)$ and

$$\mathbf{N}_m(\sigma) = \begin{bmatrix} \mathbf{I} + \sigma \mathbf{C}^* \mathbf{U} + \bar{\sigma} \mathbf{U}^* \mathbf{C} + |\sigma|^2 \mathbf{U}^* \mathbf{U} & \mathbf{B}_m + \bar{\sigma} \mathbf{U}^* \mathbf{C} \mathbf{B}_m + \bar{\sigma} \mathbf{U}^* \mathbf{V}_{m+1} \bar{\mathbf{H}}_m + |\sigma|^2 \mathbf{U}^* \mathbf{V}_m \\ \mathbf{B}_m^* + \sigma \mathbf{B}_m^* \mathbf{C}^* \mathbf{U} + \sigma \bar{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{U} + |\sigma|^2 \mathbf{V}_m^* \mathbf{U} & \mathbf{B}_m^* \mathbf{B}_m + \bar{\mathbf{H}}_m^* \bar{\mathbf{H}}_m + \sigma \mathbf{H}_m + \bar{\sigma} \mathbf{H}_m^* + |\sigma|^2 \mathbf{I} \end{bmatrix}.$$

This projection can be performed using already computed quantities, and the matrices $\mathbf{U}^* \mathbf{C}$, $\mathbf{U}^* \mathbf{U}$, $\mathbf{U}^* \mathbf{C} \mathbf{B}_m$, $\mathbf{U}^* \mathbf{V}_{m+1} \bar{\mathbf{H}}_m$, $\bar{\mathbf{H}}_m^* \bar{\mathbf{H}}_m$, \mathbf{H}_m , and $\mathbf{B}_m^* \mathbf{B}_m$ need only be computed once, regardless of the number of shifts. The only computational growth we have that is shift dependent comes from the solution of a linear system with $\mathbf{N}_m(\sigma)$ for every shift at a cost of $\mathcal{O}((m+k)^3)$.

3.2.1. Preconditioning. Introducing right preconditioning creates some difficulties which we can again surmount by storing some extra vectors. In this case, for right preconditioned Recycled GMRES, the search space for the base system is $\mathcal{S}_m \leftarrow \mathbf{M}^{-1} \{\mathcal{U} + \mathcal{K}_m(\mathbf{P}\mathbf{A}, \mathbf{r}_0)\}$, but the basis for this space does not need to be stored. However, it is available at no additional cost during the computation and can be saved. Let $\mathbf{Z}_{\mathcal{U}} = \mathbf{M}^{-1} \mathbf{U}$ and $\mathbf{Z}_m = \mathbf{M}^{-1} \mathbf{V}_m$, as in Section 3.1.

Using $\mathbf{Z}_{\mathcal{U}}$, we can cheaply perform the initial residual projection,

$$\begin{aligned} \mathbf{x}_0(\sigma) &= \mathbf{x}_{-1}(\sigma) + \mathbf{U} \mathbf{y}_0(\sigma) \quad \text{and} \\ \mathbf{r}_0(\sigma) &= \mathbf{r}_{-1}(\sigma) - (\mathbf{A}_p(\sigma) \mathbf{U}) \mathbf{y}_0(\sigma) \end{aligned} \quad (3.9)$$

where $\mathbf{y}_0(\sigma) = \mathbf{N}_0(\sigma)^{-1} (\mathbf{A}_p(\sigma) \mathbf{U})^* \mathbf{r}_{-1}(\sigma)$ and $\mathbf{N}_0(\sigma) = (\mathbf{A}_p(\sigma) \mathbf{U})^* (\mathbf{A}_p(\sigma) \mathbf{U})$. We can write

$$\mathbf{A}_p(\sigma) \mathbf{U} = \mathbf{C} + \sigma \mathbf{Z}_{\mathcal{U}}.$$

The subspace \mathcal{U} either is available from at the start of the algorithm (in which case \mathbf{U} must be scaled so that $\mathbf{A}_p \mathbf{U} = \mathbf{C}$ has orthonormal columns), or it is constructed at the end of a restart cycle. In either case, $\mathbf{Z}_{\mathcal{U}}$ is available in the course of the computation and can be saved. Thus the projection (3.7) can be performed with already computed quantities, and this computation is similar to the unpreconditioned case,

$$\begin{aligned} \mathbf{x}_0(\sigma) &= \mathbf{x}_{-1}(\sigma) + \mathbf{U} \mathbf{y}_0(\sigma) \quad \text{and} \\ \mathbf{r}_0(\sigma) &= \mathbf{r}_{-1}(\sigma) - (\mathbf{C} + \sigma \mathbf{Z}_{\mathcal{U}}) \mathbf{y}_0(\sigma), \end{aligned} \quad (3.10)$$

where we rewrite $\mathbf{y}_0(\sigma) = \mathbf{N}_0(\sigma)^{-1} (\mathbf{C} + \sigma \mathbf{Z}_{\mathcal{U}})^* \mathbf{r}_{-1}(\sigma)$ and

$$\mathbf{N}_0(\sigma) = \mathbf{I} + \sigma \mathbf{C}^* \mathbf{Z}_{\mathcal{U}} + \bar{\sigma} \mathbf{Z}_{\mathcal{U}}^* \mathbf{C} + |\sigma|^2 \mathbf{Z}_{\mathcal{U}}^* \mathbf{Z}_{\mathcal{U}}.$$

After a cycle of right-preconditioned Recycled GMRES, we must perform the projection (3.4) for each shifted system. We proceed slightly differently in this derivation than in the unpreconditioned case. We have

$$\mathbf{N}_m(\sigma) \leftarrow \{\mathbf{A}_p(\sigma) [\mathbf{U} \quad \mathbf{V}_m]\}^* \{\mathbf{A}_p(\sigma) [\mathbf{U} \quad \mathbf{V}_m]\}$$

Following from [19], we define

$$\bar{\mathbf{G}}_m = \begin{bmatrix} \mathbf{I}_{k \times k} & \mathbf{B}_m \\ \mathbf{0}_{(m+1) \times k} & \bar{\mathbf{H}}_m \end{bmatrix},$$

which yields the augmented Arnoldi relation

$$\mathbf{A}_p [\mathbf{U} \quad \mathbf{V}_m] = [\mathbf{C} \quad \mathbf{V}_{m+1}] \bar{\mathbf{G}}_m. \quad (3.11)$$

Using the relation (3.11), an identity for the shifted operator with right preconditioning follows,

$$\mathbf{A}_p(\sigma) \begin{bmatrix} \mathbf{U} & \mathbf{V}_m \end{bmatrix} = \begin{bmatrix} \mathbf{C} & \mathbf{V}_{m+1} \end{bmatrix} \overline{\mathbf{G}}_m + \sigma \begin{bmatrix} \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_m \end{bmatrix}. \quad (3.12)$$

We use the relation (3.12) to derive the expansion

$$\begin{aligned} \mathbf{N}_m(\sigma) &= \overline{\mathbf{G}}_m^* \overline{\mathbf{G}}_m + |\sigma|^2 \begin{bmatrix} \mathbf{Z}_{\mathcal{U}}^* \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_{\mathcal{U}}^* \mathbf{Z}_m \\ \mathbf{Z}_m^* \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_m^* \mathbf{Z}_m \end{bmatrix} + \sigma \overline{\mathbf{G}}_m^* \begin{bmatrix} \mathbf{C}^* \mathbf{Z}_{\mathcal{U}} & \mathbf{C}^* \mathbf{Z}_m \\ \mathbf{V}_{m+1}^* \mathbf{Z}_{\mathcal{U}} & \mathbf{V}_{m+1}^* \mathbf{Z}_m \end{bmatrix} \\ &\quad + \overline{\sigma} \begin{bmatrix} \mathbf{Z}_{\mathcal{U}}^* \mathbf{C} & \mathbf{Z}_{\mathcal{U}}^* \mathbf{V}_{m+1} \\ \mathbf{Z}_m^* \mathbf{C} & \mathbf{Z}_m^* \mathbf{V}_{m+1} \end{bmatrix} \overline{\mathbf{G}}_m. \end{aligned} \quad (3.13)$$

Thus, the projection can be performed for each shift using already computed quantities. This yields the following updates of the approximation and residual

$$\mathbf{x}_m(\sigma) = \mathbf{x}_0(\sigma) + \begin{bmatrix} \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_m \end{bmatrix} \mathbf{y}_m(\sigma) \quad (3.14)$$

$$\mathbf{r}_m(\sigma) = \mathbf{r}_0(\sigma) - \left\{ \begin{bmatrix} \mathbf{C} & \mathbf{V}_{m+1} \end{bmatrix} \overline{\mathbf{G}}_m + \sigma \begin{bmatrix} \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_m \end{bmatrix} \right\} \mathbf{y}_m(\sigma) \quad (3.15)$$

where $\mathbf{y}_m(\sigma) = \mathbf{N}_m(\sigma)^{-1} \left\{ \begin{bmatrix} \mathbf{C} & \mathbf{V}_{m+1} \end{bmatrix} \overline{\mathbf{G}}_m + \sigma \begin{bmatrix} \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_m \end{bmatrix} \right\}^* \mathbf{r}_0(\sigma)$. We observe that because of the initial projection of the shifted residual (3.9), we can simplify

$$\begin{aligned} \left\{ \begin{bmatrix} \mathbf{C} & \mathbf{V}_{m+1} \end{bmatrix} \overline{\mathbf{G}}_m + \sigma \begin{bmatrix} \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_m \end{bmatrix} \right\}^* \mathbf{r}_0(\sigma) &= \left\{ \begin{bmatrix} \mathbf{C} & \mathbf{C}\mathbf{B}_m + \mathbf{V}_{m+1}\overline{\mathbf{H}}_m \end{bmatrix} + \sigma \begin{bmatrix} \mathbf{Z}_{\mathcal{U}} & \mathbf{Z}_m \end{bmatrix} \right\}^* \mathbf{r}_0(\sigma) \\ &= \begin{bmatrix} \mathbf{C}^* \mathbf{r}_0(\sigma) \\ \mathbf{B}_m^* \mathbf{C}^* \mathbf{r}_0(\sigma) + \overline{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{r}_0(\sigma) \end{bmatrix} + \overline{\sigma} \begin{bmatrix} \mathbf{Z}_{\mathcal{U}}^* \mathbf{r}_0(\sigma) \\ \mathbf{Z}_m^* \mathbf{r}_0(\sigma) \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_m^* \mathbf{C}^* \mathbf{r}_0(\sigma) + \overline{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{r}_0(\sigma) + \overline{\sigma} \mathbf{Z}_m^* \mathbf{r}_0(\sigma) \end{bmatrix}, \end{aligned}$$

and thus we can rewrite

$$\mathbf{y}_m(\sigma) = \mathbf{N}_m(\sigma)^{-1} \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_m^* \mathbf{C}^* \mathbf{r}_0(\sigma) + \overline{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{r}_0(\sigma) + \overline{\sigma} \mathbf{Z}_m^* \mathbf{r}_0(\sigma) \end{bmatrix}.$$

The matrices in the sum (3.13) must be computed only once. As before, for each shift, a linear system with coefficient matrix $\mathbf{N}_m(\sigma)$ must be solved at a cost of $\mathcal{O}((m+k)^3)$. The right-preconditioned shifted Recycled GMRES algorithm (sometimes referred to as srGMRES) is shown in Algorithm 3.2. Observe that an implementation can rely heavily on an existing Recycled GMRES code. As in the case of Algorithm 3.1, all but one step of the shifted residual projections can be formulated in terms of block/BLAS-3 operations so that the computations are performed simultaneously for all shifts. However, we again must compute $\mathbf{y}_m(\sigma)$ for each shift separately, as this computation involves the inverse of $\mathbf{N}_m(\sigma)$.

4. Analysis of direct projection methods. In this section, we provide some analysis of the direction projection methods. For ease of notation and to provide analysis which applies to all methods which fit into this framework including those developed in Section 3, this analysis is developed in the general framework setting and then interpreted for the individual methods.

Before going further, we note that since all residual corrections are minimum residual projections, we can expect that, at worse, the projection of the shifted residual will achieve no improvement. This procedure will not increase the shifted residual norm.

Our analysis proceeds in two distinct directions:

- recognizing that they are related to techniques developed in the 1980's for symmetric linear systems with multiple right-hand sides.

Algorithm 3.2: Right preconditioned shifted Recycled GMRES (**srGMRES()**)

Input : $\mathbf{A} \in \mathbb{C}^{n \times n}$; $\mathbf{b} \in \mathbb{C}^n$; $\{\sigma_\ell\}_{\ell=1}^L \subset \mathbb{C}$; Initial Approximations $\{\mathbf{x}(\sigma_\ell)\}_{\ell=1}^L$;
 $\mathbf{U} \in \mathbb{C}^{n \times k}$; $\varepsilon > 0$; Cycle length $m \in \mathbb{N}$

Output: $\{\mathbf{x}(\sigma_\ell)\}_{\ell=1}^L$ such that $\|\mathbf{r}(\sigma_\ell)\| / \|\mathbf{r}_0(\sigma_\ell)\| \leq \varepsilon$ for all ℓ

```

1 for  $\ell = 1 \dots L$  do
2    $\mathbf{r}(\sigma_\ell) = \mathbf{b} - (\mathbf{A} + \sigma_\ell \mathbf{I})\mathbf{x}(\sigma_\ell)$ 
3    $\gamma_1 = \|\mathbf{r}(\sigma_1)\|$ 
4    $\mathbf{Z}_\mathcal{U} = \mathbf{M}^{-1}\mathbf{U}$ 
5    $\mathbf{C} = (\mathbf{A} + \sigma_1 \mathbf{I})\mathbf{Z}_\mathcal{U}$ 
6   Compute QR-factorization  $\mathbf{QR} = \mathbf{C}$ 
7    $\mathbf{C} \leftarrow \mathbf{Q}$ ,  $\mathbf{U} \leftarrow \mathbf{UR}^{-1}$ 
8    $\mathbf{x}(\sigma_1) \leftarrow \mathbf{x}(\sigma_1) + \mathbf{UC}^*\mathbf{r}(\sigma_1)$  and  $\mathbf{r}(\sigma_1) \leftarrow \mathbf{r}(\sigma_1) - \mathbf{CC}^*\mathbf{r}(\sigma_1)$ 
9   Compute  $\mathbf{C}^*\mathbf{Z}_\mathcal{U}$  and  $\mathbf{Z}_\mathcal{U}^*\mathbf{Z}_\mathcal{U}$ 
10  for  $\ell = 2 \dots L$  do
11    %%%% Shifted System Initial Projections %%%%
12     $\mathbf{N} \leftarrow \mathbf{I} + \sigma \mathbf{C}^*\mathbf{Z}_\mathcal{U} + \bar{\sigma} \mathbf{Z}_\mathcal{U}^* \mathbf{C} + |\sigma|^2 \mathbf{Z}_\mathcal{U}^* \mathbf{Z}_\mathcal{U}$ 
13     $\mathbf{y} \leftarrow \mathbf{N}^{-1} (\mathbf{C} + \sigma \mathbf{Z}_\mathcal{U})^* \mathbf{r}(\sigma_\ell)$ 
14     $\mathbf{x}(\sigma_\ell) \leftarrow \mathbf{x}(\sigma_\ell) + \mathbf{Uy}$ 
15     $\mathbf{r}(\sigma_\ell) \leftarrow \mathbf{r}(\sigma_\ell) - (\mathbf{C} + \sigma \mathbf{Z}_\mathcal{U}) \mathbf{y}$ 
16  if  $L > 1$  then
17    while  $\|\mathbf{r}(\sigma_1)\| / \gamma_1 > \varepsilon$  do
18      Compute and overwrite  $\mathbf{x}(\sigma_1)$ ,  $\mathbf{r}(\sigma_1)$ ,  $\mathbf{V}_{m+1}$ ,  $\mathbf{Z}_m$ ,  $\bar{\mathbf{H}}_m$ ,  $\mathbf{B}_m$  by calling
19      rGMRES() for  $\mathbf{A} + \sigma_1 \mathbf{I}$ ,  $\mathbf{M}$ ,  $\mathbf{b}$ ,  $\mathbf{x}(\sigma_1)$ ,  $\mathbf{U}$ ,  $\mathbf{C}$ , and  $m$ 
20      Compute and overwrite  $\bar{\mathbf{G}}_m$ ,  $\bar{\mathbf{G}}_m^* \bar{\mathbf{G}}_m$ ,  $\mathbf{Z}_\mathcal{U}^* \mathbf{Z}_\mathcal{U}$ ,  $\mathbf{Z}_\mathcal{U}^* \mathbf{Z}_m$ ,  $\mathbf{Z}_m^* \mathbf{Z}_m$ ,  $\mathbf{C}^* \mathbf{Z}_\mathcal{U}$ ,
21       $\mathbf{C}^* \mathbf{Z}_m$ ,  $\mathbf{V}_{m+1}^* \mathbf{Z}_\mathcal{U}$ ,  $\mathbf{V}_{m+1}^* \mathbf{Z}_m$ 
22      for  $\ell = 2 \dots s$  do
23        %%%% Shifted System Projections %%%%
24         $\mathbf{N} \leftarrow \bar{\mathbf{G}}_m^* \bar{\mathbf{G}}_m + |\sigma|^2 \begin{bmatrix} \mathbf{Z}_\mathcal{U}^* \mathbf{Z}_\mathcal{U} & \mathbf{Z}_\mathcal{U}^* \mathbf{Z}_m \\ \mathbf{Z}_m^* \mathbf{Z}_\mathcal{U} & \mathbf{Z}_m^* \mathbf{Z}_m \end{bmatrix} +$ 
25         $\sigma \bar{\mathbf{G}}_m^* \begin{bmatrix} \mathbf{C}^* \mathbf{Z}_\mathcal{U} & \mathbf{C}^* \mathbf{Z}_m \\ \mathbf{V}_{m+1}^* \mathbf{Z}_\mathcal{U} & \mathbf{V}_{m+1}^* \mathbf{Z}_m \end{bmatrix} + \bar{\sigma} \begin{bmatrix} \mathbf{Z}_\mathcal{U}^* \mathbf{C} & \mathbf{Z}_\mathcal{U}^* \mathbf{V}_{m+1} \\ \mathbf{Z}_m^* \mathbf{C} & \mathbf{Z}_m^* \mathbf{V}_{m+1} \end{bmatrix} \bar{\mathbf{G}}_m$ 
26         $\mathbf{y} \leftarrow \mathbf{N}^{-1} \{ [\mathbf{C} \quad \mathbf{V}_{m+1}] \bar{\mathbf{G}}_m + \sigma [\mathbf{Z}_\mathcal{U} \quad \mathbf{Z}_m] \}^* \mathbf{r}(\sigma_\ell)$ 
27         $\mathbf{x}(\sigma_\ell) \leftarrow \mathbf{x}_0(\sigma_\ell) + [\mathbf{Z}_\mathcal{U} \quad \mathbf{Z}_m] \mathbf{y}$ 
28         $\mathbf{r}(\sigma_\ell) \leftarrow \mathbf{r}_0(\sigma_\ell) - \{ [\mathbf{C} \quad \mathbf{V}_{m+1}] \bar{\mathbf{G}}_m + \sigma [\mathbf{Z}_\mathcal{U} \quad \mathbf{Z}_m] \} \mathbf{y}$ 
29      Compute updated  $\mathbf{U}$ ,  $\mathbf{Z}_\mathcal{U}$ , and  $\mathbf{C}$ 
30    For all  $\ell = 2, \dots, L$  compute and overwrite  $\mathbf{x}(\sigma_\ell)$  by recursively calling
31    srGMRES() for  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $\mathbf{M}$ ,  $\{\sigma_\ell\}_{\ell=2}^L$ ,  $\{\mathbf{x}(\sigma_\ell)\}_{\ell=2}^L$ ,  $\mathbf{U}$ ,  $\varepsilon$ , and  $m$ 
32  else
33    while  $\|\mathbf{r}(\sigma_1)\| / \gamma_1 > \varepsilon$  do
34      Compute and overwrite  $\mathbf{x}(\sigma_1)$ ,  $\mathbf{r}(\sigma_1)$  by calling rGMRES() for  $\mathbf{A} + \sigma_1 \mathbf{I}$ ,
35       $\mathbf{M}$ ,  $\mathbf{b}$ ,  $\mathbf{x}(\sigma_1)$ , and  $m$ 

```

- studying the projection scaling matrix $\mathbf{N}_m(\sigma)$, showing that it can be related to an eigenvalue approximation problem associated to the search space.

Direct applying a projection of this type to each shifted system is an application (in the non-Hermitian case) of the Lanczos-Galerkin technique; see, e.g., [4, 20, 21]. Using analytic tools developed for such methods, we can derive a bound on the residual norm for a shifted system after projection.

Following from [21], we analyze the direct projection methods by decomposing the residual in a particular way, described in the following lemma.

LEMMA 4.1. *Let the sequence of subspaces $\{\mathcal{S}_m\}$ be defined as in Section 3, and additionally let*

$$\mathcal{T}_m = \mathbf{A}\mathcal{S}_m. \quad (4.1)$$

Let $\overline{\mathcal{T}}_m = \mathcal{T}_m + \sigma\mathcal{S}_{m-1}$, and define \mathbf{P}_m to be the orthogonal projector onto $(\mathbf{A} + \sigma\mathbf{I})\mathcal{S}_m$ and $\mathbf{P}_{\overline{\mathcal{T}}_m}$ to be the orthogonal projector onto $\overline{\mathcal{T}}_m$. If $\mathbf{r}_0(\sigma)$ is the initial residual for the shifted system, and $\mathbf{r}_m(\sigma)$ is the residual produced by projecting $\mathbf{r}_0(\sigma)$ according to (3.4), then we have that

$$\mathbf{r}_m(\sigma) = (\mathbf{I} - \mathbf{P}_m)\mathbf{P}_{\overline{\mathcal{T}}_{m+1}}\mathbf{r}_0(\sigma) + (\mathbf{I} - \mathbf{P}_{\overline{\mathcal{T}}_m})\mathbf{r}_0(\sigma). \quad (4.2)$$

Proof. Using the property of projectors, we can decompose

$$\mathbf{r}_0(\sigma) = \mathbf{P}_{\overline{\mathcal{T}}_m}\mathbf{r}_0(\sigma) + (\mathbf{I} - \mathbf{P}_{\overline{\mathcal{T}}_m})\mathbf{r}_0(\sigma).$$

The minimum residual projection (3.4) can be written,

$$\mathbf{r}_m(\sigma) = (\mathbf{I} - \mathbf{P}_m)\mathbf{P}_{\overline{\mathcal{T}}_{m+1}}\mathbf{r}_0(\sigma) + (\mathbf{I} - \mathbf{P}_m)(\mathbf{I} - \mathbf{P}_{\overline{\mathcal{T}}_{m+1}})\mathbf{r}_0(\sigma). \quad (4.3)$$

From (4.1) and the definition of $\overline{\mathcal{T}}_m$, we have that

$$(\mathbf{A} + \sigma\mathbf{I})\mathcal{S}_m \subset \overline{\mathcal{T}}_{m+1}$$

which in turn yields the reverse containment of the orthogonal complements,

$$\overline{\mathcal{T}}_{m+1}^\perp \subset \{(\mathbf{A} + \sigma\mathbf{I})\mathcal{T}_m\}^\perp$$

and thus

$$(\mathbf{I} - \mathbf{P}_m)(\mathbf{I} - \mathbf{P}_{\overline{\mathcal{T}}_{m+1}})\mathbf{r}_0(\sigma) = (\mathbf{I} - \mathbf{P}_{\overline{\mathcal{T}}_{m+1}})\mathbf{r}_0(\sigma)$$

This yields the result. \square

COROLLARY 4.2. *Let the same assumptions as in Lemma 4.1 hold. Then we have the following bound on $\|\mathbf{r}_m(\sigma)\|$,*

$$\|\mathbf{r}_m(\sigma)\| \leq \|(\mathbf{I} - \mathbf{P}_m)\mathbf{P}_{\overline{\mathcal{T}}_{m+1}}\mathbf{r}_0(\sigma)\| + \|(\mathbf{I} - \mathbf{P}_{\overline{\mathcal{T}}_m})\mathbf{r}_0(\sigma)\| \quad (4.4)$$

Proof. We simply take the norm of both sides of (4.2) and apply the triangle inequality. \square

From (4.4), we can see that the residual norm bound depends on both the effectiveness of the Petrov-Galerkin minimization projection applied to the orthogonal projection of $\mathbf{r}_0(\sigma)$ in $\overline{\mathcal{T}}_{m+1}$ and the size of the part of the residual which lies in $\overline{\mathcal{T}}_{m+1}^\perp$. As an aside, to connect this analysis back to the two proposed methods, we observe

that in the case of the right-preconditioned shifted GMRES algorithm (Algorithm 3.1), we have

$$\begin{aligned}\mathcal{S}_m &= \mathbf{M}^{-1}\mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0), \quad \mathcal{T}_m = \mathbf{A}\mathbf{M}^{-1}\mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0), \quad \text{and,} \\ \overline{\mathcal{T}}_m &= \mathbf{A}\mathbf{M}^{-1}\mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0) + \sigma\mathbf{M}^{-1}\mathcal{K}_m(\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0).\end{aligned}\quad (4.5)$$

In the case of the preconditioned rGMRES method for shifted systems (Algorithm 3.2), we have

$$\begin{aligned}\mathcal{S}_m &= \mathbf{M}^{-1}\{\mathcal{U} + \mathcal{K}_m((\mathbf{I} - \mathbf{P})\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0)\}, \quad \mathcal{T}_m = \mathcal{C} + \mathbf{A}\mathbf{M}^{-1}\mathcal{K}_m((\mathbf{I} - \mathbf{P})\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0), \quad \text{and,} \\ \overline{\mathcal{T}}_m &= \mathcal{C} + \mathbf{A}\mathbf{M}^{-1}\mathcal{K}_m((\mathbf{I} - \mathbf{P})\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0) + \sigma\mathbf{M}^{-1}\{\mathcal{U} + \mathcal{K}_m((\mathbf{I} - \mathbf{P})\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0)\}\end{aligned}\quad (4.6)$$

We turn now to the matrix $\mathbf{N}_m(\sigma)$, and show that it is connected to an eigenvalue approximation problem associated to the search space (i.e., computation of the harmonic Ritz values [17, 18]). Recall that for a matrix \mathbf{A} and a subspace \mathcal{S} , a pair $(\mathbf{u}, 1/\mu)$ with $\mathbf{u} \in \mathbb{R}^n$ and $\mu \in \mathbb{C}$ is called a *harmonic Ritz pair* (harmonic Ritz value and vector individually) with respect to \mathbf{A} and \mathcal{S} if $\mathbf{u} \in \mathbf{A}\mathcal{S}$ and the pair satisfies the orthogonality condition

$$\mathbf{A}^{-1}\mathbf{u} - \mu\mathbf{u} \perp \mathbf{A}\mathcal{S}. \quad (4.7)$$

When $\mathcal{S} = \mathcal{K}_m(\mathbf{A}, \mathbf{r}_0)$ is a Krylov subspace, the harmonic Ritz values are closely related to the GMRES minimal residual solution associated to that subspace. The values themselves are roots of the GMRES minimum residual polynomial [12], and they can be computed as the eigenvalues of the matrix $(\overline{\mathbf{H}}_m^* \overline{\mathbf{H}}_m)^{-1} \mathbf{H}_m$ [17], where $\mathbf{H}_m \in \mathbb{C}^{m \times m}$ contains the first m rows of $\overline{\mathbf{H}}_m$. This matrix is also the system matrix from the normal equations of the GMRES least squares problem (2.4). This can be seen by rewriting the normal equations associated to (2.4) which are

$$\overline{\mathbf{H}}_m^* \overline{\mathbf{H}}_m \mathbf{y} = \|\mathbf{r}_0\| \overline{\mathbf{H}}_m^* \mathbf{e}_1^{j+1}. \quad (4.8)$$

We observe here that $\overline{\mathbf{H}}_m^* \mathbf{e}_1^{j+1}$ is simply the first column of $\overline{\mathbf{H}}_m^*$, which is also the first column of \mathbf{H}_m^* . Thus, we can rewrite (4.8)

$$\overline{\mathbf{H}}_m^* \overline{\mathbf{H}}_m \mathbf{y} = \|\mathbf{r}_0\| \mathbf{H}_m^* \mathbf{e}_1^j. \quad (4.9)$$

Furthermore, we can write explicitly the GMRES projection

$$\begin{aligned}\mathbf{r}_m &= \mathbf{r}_0 - \mathbf{A}\mathbf{V}_m((\mathbf{A}\mathbf{V}_m)^* \mathbf{A}\mathbf{V}_m)^{-1}(\mathbf{A}\mathbf{V}_m)^* \mathbf{r}_0 \\ &= \mathbf{r}_0 - \mathbf{V}_{m+1} \overline{\mathbf{H}}_m (\overline{\mathbf{H}}_m^* \overline{\mathbf{H}}_m)^{-1} \overline{\mathbf{H}}_m^* \mathbf{V}_{m+1}^* \mathbf{r}_0 \\ &= \mathbf{r}_0 - \mathbf{V}_{m+1} \overline{\mathbf{H}}_m (\overline{\mathbf{H}}_m^* \overline{\mathbf{H}}_m)^{-1} \overline{\mathbf{H}}_m^* \|\mathbf{r}_0\| \mathbf{e}_1^{j+1} \\ &= \mathbf{r}_0 - \mathbf{V}_{m+1} \overline{\mathbf{H}}_m (\overline{\mathbf{H}}_m^* \overline{\mathbf{H}}_m)^{-1} \mathbf{H}_m^* \|\mathbf{r}_0\| \mathbf{e}_1^j,\end{aligned}\quad (4.10)$$

and we see the same matrix appear. We see here that the GMRES minimum residual projection scaling matrix is $\overline{\mathbf{H}}_m^* \overline{\mathbf{H}}_m$, the matrix whose inverse appears in the harmonic Ritz matrix as well as in the solution of the GMRES least squares problem.

Can we derive similar relations in our general framework? Certainly, there is no underlying residual polynomial interpolation problem, but we can still compute harmonic Ritz values for the search space \mathcal{S} in our framework.

PROPOSITION 4.3. *Each harmonic Ritz pair with respect to $\mathbf{A}(\sigma)$ and \mathcal{S}_m can be written as $(\mathbf{A}(\sigma)\mathbf{S}_m\mathbf{y}, 1/\mu)$, where (\mathbf{y}, μ) is some eigenpair from the following generalized eigenvalue problem,*

$$(\mathbf{A}(\sigma)\mathbf{S}_m)^*\mathbf{S}_m\mathbf{y} = \mu\mathbf{N}_\sigma\mathbf{y}.$$

Proof. From (4.7), at iteration m , we seek $u = \mathbf{A}(\sigma)\mathbf{S}_m\mathbf{y}$ and μ such that

$$\mathbf{A}(\sigma)^{-1}\mathbf{u} - \mu\mathbf{u} \perp \mathbf{A}(\sigma)\mathcal{S}_m \quad (4.11)$$

which we can also write as the equation

$$\begin{aligned} (\mathbf{A}(\sigma)\mathbf{S}_m)^* (\mathbf{A}(\sigma)^{-1}\mathbf{u} - \mu\mathbf{u}) &= 0 \\ (\mathbf{A}(\sigma)\mathbf{S}_m)^* (\mathbf{A}(\sigma)^{-1}\mathbf{A}(\sigma)\mathbf{S}_m\mathbf{y} - \mu\mathbf{A}(\sigma)\mathbf{S}_m\mathbf{y}) &= 0 \\ (\mathbf{A}(\sigma)\mathbf{S}_m)^* (\mathbf{S}_m\mathbf{y} - \mu\mathbf{A}(\sigma)\mathbf{S}_m\mathbf{y}) &= 0 \\ (\mathbf{A}(\sigma)\mathbf{S}_m)^*\mathbf{S}_m\mathbf{y} &= \mu\mathbf{N}_{m(\sigma)}\mathbf{y}, \end{aligned}$$

yielding the result. \square

Thus we see that as in the GMRES case, the inverse of our projection scaling matrix $\mathbf{N}_{m(\sigma)}$ appears as a factor in the resulting generalized eigenvalue problem.

We can also recover the minimum residual correction in \mathcal{S}_m (which is computed directly in the projector) as the solution of a small linear system by following from the general description of projection methods described in [22, Section 5.2.3]. However, we make this observation and omit the redundant computations for the sake of brevity.

We comment that the preceding analysis does not paint a full picture of the expected behavior of the proposed methods. As we will see in numerical experiments, the amount of residual improvement derived from the shifted system projections degrade as the magnitude of the shifts increases, though this was not consistently observed, c.f. Figure 5.4. However, we do observe that, in these experiments, the resulting residuals do seem to produce a quickly converging iteration when the systems to which they are associated become the base systems upon a recursive function call. Our analysis does not yet pinpoint the cause of this phenomenon or describe it quantitatively.

5. Numerical Results. We performed a series of numerical experiments to demonstrate the effectiveness of our algorithms as well as to compare performance (as measured in both matrix-vector product counts and CPU timings) with other algorithms. All tests were performed in Matlab R2013a (8.1.0.604) running on a Macbook Pro with a 2.3 GHz Intel Core i5 processor and 8 GB of 1333 MHz DDR3 main memory. For these tests, we use two sets of QCD matrices downloaded from the University of Florida Sparse Matrix Library [6]. One set of matrices is a collection of seven 3072×3072 complex matrices and the other is a collection of seven 49152×49152 complex matrices. For all experiments, we chose the right-hand side $\mathbf{b} = \mathbf{1}$, the vector of ones and the requested relative residual tolerance $\varepsilon = 10^{-8}$. For all experiments, we preconditioned with an incomplete LU-factorization (ILU) constructed using the Matlab function `ilu()` called with the default Matlab settings. We comment that the usage of ILU was a matter of convenience and effectiveness for these sample problems. Its usage is meant to demonstrate proof-of-concept rather than as advocating the usage of ILU for large-scale QCD problems.

We also comment about two methods which we have omitted from testing, the shifted restarted GMRES method [8] and the recursive Recycled GMRES method for shifted systems proposed in [28]. We have omitted these methods from the tests as

they do not admit general preconditioning. As such, they require substantially more iterations in many experiments. However, with the method of Frommer and Glässner, there would be some number of shifts for which this method would be superior to those presented in this paper, as cost of recursion in our methods, even with preconditioning, would be greater than simply solving the unpreconditioned problems simultaneously with their shifted GMRES method [8].

Since these experiments involve solving shifted systems with shifts of varying magnitudes, it is useful to know information about the norms of our test matrices. Therefore, we provide both the one- and two-norms for these matrices (computed respectively with the Matlab functions `norm(·, 1)` and `svds(·, 1)`) in Tables 5.1 and 5.2.

TABLE 5.1
Norms of Small 3072×3072 QCD Matrices

	$\mathbf{A}_1^{(s)}$	$\mathbf{A}_2^{(s)}$	$\mathbf{A}_3^{(s)}$	$\mathbf{A}_4^{(s)}$	$\mathbf{A}_5^{(s)}$	$\mathbf{A}_6^{(s)}$	$\mathbf{A}_7^{(s)}$
$\ \mathbf{A}_i\ _1^{(s)}$	28.734	28.763	28.705	28.746	28.632	29.931	29.732
$\ \mathbf{A}_i\ _2^{(s)}$	11.272	11.294	11.37	11.332	11.054	13.211	12.864

TABLE 5.2
Norms of Large 49752×49752 QCD Matrices

	$\mathbf{A}_1^{(L)}$	$\mathbf{A}_2^{(L)}$	$\mathbf{A}_3^{(L)}$	$\mathbf{A}_4^{(L)}$	$\mathbf{A}_5^{(L)}$	$\mathbf{A}_6^{(L)}$	$\mathbf{A}_7^{(L)}$
$\ \mathbf{A}_i^{(L)}\ _1$	29.509	29.576	29.568	29.504	30.324	30.305	30.397
$\ \mathbf{A}_i^{(L)}\ _2$	12.367	12.385	12.444	12.377	13.428	13.49	13.48

In our first experiment, we solved a sequence of families of shifted linear systems to show a sample convergence history for Algorithm 3.2 for the large QCD matrices, see Figure 5.1.

We then tested Algorithm 3.2 with the smaller set of matrices for various recycle space dimension sizes and restart cycle lengths. We solve for shifts $\sigma \in \{.01, .02, .03, 1, 2, 3\}$. We calculated total required matrix-vector products. In Table 5.3, we show the results. We see in the table that for these particular QCD matrices, good results can be achieved for a small recycled subspace dimension as long as the cycle length is sufficiently long.

TABLE 5.3
Matrix vector product counts for different pairs (m, k) of restart cycle length and recycled subspace dimension for shifted Recycled GMRES. Experiments were performed for larger values than shown but no further improvement was observed

$m \backslash k$	5	15	25	35	45	55	65	75	85
5	1566	1295	1205	1161	1146	1131	1126	1116	1111
20	1466	1254	1182	1141	1122	1110	1107	1103	1096
35	1418	1229	1166	1132	1113	1103	1096	1095	1091
50	1363	1223	1158	1128	1114	1105	1099	1097	1090
65	1344	1219	1159	1124	1109	1106	1099	1090	1086
80	1321	1210	1153	1123	1109	1102	1098	1091	1085
95	1321	1210	1153	1124	1108	1100	1097	1093	1084




FIG. 5.1. Convergence history for the large QCD matrices for Algorithm 3.2 with shifts $\{.01, .02, .03, 1, 2, 3\}$ and $(m, k) = (100, 5)$. In each subplot, we display the residual curves sequentially to reflect that the algorithm is called for each shifted linear system in sequence. The order in which the systems were solved is the same as the order of the listed shifts.

For the remaining tests, we use the larger set of QCD matrices. In Table 5.4 we compare time and matrix-vector product counts. Based on results from Table 5.3, we chose $(m, k) = (100, 5)$ and use this pair for all experiments except for the one shown in Figure 5.2. We choose a smaller recycled subspace dimension and a longer cycle length, in order to maintain lower per-iteration costs (the subspace recycling itself naturally has some overhead) For each family of linear systems, the experiment was performed ten times and the average time over these ten runs was taken as the run time. We solved for a larger number of shifts of varying magnitudes,

$$\sigma \in \{.001, .002, .003, .04, .05, .06, .07, .8, .9, 1, 1.1, 10, 11, 12\}.$$

We compared four methods (Algorithm 3.1, Algorithm 3.2, sequentially applied GMRES and sequentially applied Recycled GMRES). We see that for this problem with these shifts, both proposed algorithms outperform the sequential applications of GMRES and rGMRES both in terms of matrix-vector product counts and run times. In this case, the shifted GMRES algorithm is superior in time to shifted rGMRES but not in terms of matrix-vector products. We speculate this is due to the recycling overhead of the shifted rGMRES algorithm being more costly than those of shifted GMRES, as we have implemented them in MATLAB.

In Figure 5.2, for a total fixed augmented subspace dimension of 100, we investigate how many matrix vector products are required to solve the same sequence of problems with the same shifts as in the previous experiment for different values of (m, k) such that $m + k = 100$ where m is the dimension of the projected Krylov subspace and k is the dimension of the recycled subspace. With this we demonstrate a reduction in iterations as we allow more information to be retained in the subspace.

TABLE 5.4

Timing (in seconds) and matrix-vector product (mat-vec) comparisons between shifted rGMRES, shifted GMRES, and sequential applications of rGMRES with cycle length $m = 100$ and recycled subspace dimension $k = 5$.

Method	mat-vecs	time
srGMRES	3345	252.75
sGMRES	3966	249.04
Seq. rGMRES	4940	274.37
Seq. GMRES	5685	374.3




FIG. 5.2. Matrix-vector product counts for shifted Recycled GMRES for various pairs (m, k) of Krylov subspace dimension and recycled subspace dimension such that the total augmented subspace Krylov subspace dimension $m + k = 100$.

In Table 5.5, we study matrix-vector product counts for different methods for shifts of varying magnitudes. For each shift, we solve just two systems, the base system and one shifted system. Thus we can see how many additional matrix-vector products are required for shifts of different magnitudes. What we see is that for this set of matrices, overall performance does not depend on shift magnitude. For larger shifts, we see that Algorithm 3.2 and sequentially applied rGMRES are comparable **when there is only one shift**.

TABLE 5.5

Comparison of 3 methods for different shifts sizes. In each experiment, two systems were solve, the base system and one shifted system with the shift shown in the table column header.

Method \ $\ \sigma\ $	10^{-3}	10^{-2}	10^{-1}	10^0	10^1	10^2	10^3
Sh. GMRES Alg. 3.1	1171	1233	1160	847	947	1145	1184
Sh. rGMRES Alg. 3.2	1025	1087	1056	811	913	1109	1149
Seq. rGMRES	1238	1226	1116	818	920	1118	1157

However, we have seen for larger numbers of shifts that Algorithm 3.2 exhibits superior performance. This raises the question, what are the marginal costs of solving each additional linear system for Recycled GMRES and shifted Recycled GMRES, i.e., how many more matrix-vector products does each new shifted system require? For two sets of twenty shifts, we calculated the marginal cost of solving each additional shifted system using Algorithm 3.2 as compared to Recycled GMRES. The first set

of shifts (left-hand figure) were evenly space points from the interval $[0, 1]$, and the second set of shifts (right-hand figure) were evenly spaced points from the larger interval $[1, 10]$. In Figure 5.3, we see that for the smaller interval, the cost of each




FIG. 5.3. Comparison of the marginal cost of solving each addition shifted system. For the left-hand figure, the shifts were evenly space points from the interval $[0, 1]$, and in the right-hand figure, the shifts were evenly spaced points from the larger interval $[1, 10]$

new shifted system drops for both algorithms but that Algorithm 3.2 has the lower marginal cost per shift. For the larger set of shifts, we see that the marginal costs for both algorithms actually increases for each new shift. However, the marginal cost of each new shifted system for Algorithm 3.2 becomes more stable (it levels off). For sequentially applied Recycled GMRES, the marginal costs increases steadily for all twenty shifts.

In Figure 5.4, we show the residual histories for systems solved using Algorithm 3.2 for shifts of various magnitudes,

$$\sigma \in \{10^{-3}, 10^{-2}, 10^{-1}, 1, 10^1, 10^2, 10^3\}.$$

When viewing Figure 5.4, we see (in this example) that the amount of improvement for the shifted residuals is somewhat predicted by the shift magnitude, though we observe that once shift magnitude reaches roughly 10^2 , we again see improvement in the shifted residuals due to the projection, with the projection being less effective for residuals associated to shifted systems with larger shifts.

Omitted here is a study of the eigendecomposition of the residuals, which yielded no discernible damping of certain eigenmodes or other interesting observable phenomena after the projection of the shifted residuals.

6. Conclusions. We have presented two new methods for solving a family (or a sequence of families) of shifted linear systems with general preconditioning. These methods are derived from a general framework, which we also developed in this paper. These methods use subspaces generated during the minimum residual iteration of the base system to inexpensively perform the projections for the shifted systems. This technique is fully compatible with right preconditioning, requiring only some additional storage. The strength of methods derived from this framework is that preconditioned methods for shifted systems easily can be built on top of existing minimum residual projection algorithms (and existing codes) with only minor modifications. We developed two algorithms: shifted GMRES and shifted Recycled GMRES. We




FIG. 5.4. For the large QCD matrices and $(m, k) = (100, 5)$, an illustration of the amount of residual improvement for different magnitude shifts, $\sigma \in \{10^{-3}, 10^{-2}, \dots, 10^3\}$. In each subplot, we display the residual curves sequentially to reflect that the algorithm is called for each shifted linear system in sequence. The order in which the systems were solved is the same as the order of the listed shifts.

demonstrated with numerical experiments that both methods can perform competitively.

Finally, we note that our framework is fully compatible with flexible and inexact Krylov subspace methods.

Acknowledgments. The author would like to thank Michael Parks who, while reviewing the author's dissertation, made a comment which inspired this work. The author would also like to thank Valeria Simoncini for insightful questions and comments during the author's visit to Bologna and Daniel Szyld for constructive comments.

REFERENCES

- [1] MIAN ILYAS AHMAD, DANIEL B. SZYLD, AND MARTIN B. VAN GIJZEN, *Preconditioned multishift BiCG for \mathcal{H}_2 -optimal model reduction*, Tech. Report 12-06-15, Department of Mathematics, Temple University, June 2012.
- [2] ALLISON H. BAKER, ELIZABETH R. JESSUP, AND THOMAS MANTEUFFEL, *A technique for accelerating the convergence of restarted GMRES*, SIAM Journal on Matrix Analysis and Applications, 26 (2005), pp. 962–984.
- [3] MARIA R. CELIS, JOHN E. DENNIS, AND RICHARD A. TAPIA, *A trust region strategy for nonlinear equality constrained optimization*, in Numerical optimization, 1984 (Boulder, Colo., 1984), Paul T. Boggs, Richard H. Byrd, and Robert B. Schnabel, eds., SIAM, Philadelphia, PA, 1985, pp. 71–82.
- [4] TONY F. CHAN AND WING L. WAN, *Analysis of projection methods for solving linear systems with multiple right-hand sides*, SIAM Journal on Scientific Computing, 18 (1997), pp. 1698–1721.

- [5] DEAN DARNELL, RONALD B. MORGAN, AND WALTER WILCOX, *Deflated GMRES for systems with multiple shifts and multiple right-hand sides*, Linear Algebra and its Applications, 429 (2008), pp. 2415–2434.
- [6] TIMOTHY A. DAVIS AND YIFAN HU, *The University of Florida sparse matrix collection*, ACM Trans. Math. Softw., 38 (2011), pp. 1:1–1:25.
- [7] ANDREAS FROMMER, *BiCGStab(l) for families of shifted linear systems*, Computing, 70 (2003), pp. 87–109.
- [8] ANDREAS FROMMER AND UWE GLÄSSNER, *Restarted GMRES for shifted linear systems*, SIAM Journal on Scientific Computing, 19 (1998), pp. 15–26.
- [9] ANDREAS FROMMER, STEPHAN GÜSKEN, THOMAS LIPPERT, BERTOLD NÖCKEL, AND KLAUS SCHILLING, *Many masses on one stroke: Economic computation of quark propagators*, International Journal of Modern Physics C, 6 (1995), pp. 627–638.
- [10] ANDRÉ GAUL, MARTIN H. GUTKNECHT, JÖRG LIESEN, AND REINHARD NABBEN, *A framework for deflated and augmented Krylov subspace methods*, SIAM Journal on Matrix Analysis and Applications, 34 (2013), pp. 495–518.
- [11] ANDRÉ GAUL AND NICO SCHLÖMER, *Preconditioned recycling krylov subspace methods for self-adjoint problems*, eprint 1208.0264, arXiv, 2013.
- [12] SERGE GOOSSENS AND DIRK ROOSE, *Ritz and harmonic Ritz values and the convergence of FOM and GMRES*, Numerical Linear Algebra with Applications, 6 (1999), pp. 281–293.
- [13] BEAT JEGERLEHNER, *Krylov space solvers for sparse linear systems.*, Tech. Report IUHET-353, Indiana University, 1996.
- [14] SABRINA KIRCHNER, *IDR-Verfahren zur Lösung von Familien gestifteter linearer Gleichungssysteme*, master's thesis, Bergische Universität Wuppertal, Department of Mathematics, Wuppertal, Germany, 2011.
- [15] RICHARD B. LEHOUCQ AND DANNY C. SORENSEN, *Deflation techniques for an implicitly restarted Arnoldi iteration*, SIAM Journal on Matrix Analysis and Applications, 17 (1996), pp. 789–821.
- [16] RONALD B. MORGAN, *GMRES with deflated restarting*, SIAM Journal on Scientific Computing, 24 (2002), pp. 20–37.
- [17] RONALD B. MORGAN AND MIN ZENG, *Harmonic projection methods for large non-symmetric eigenvalue problems*, Numerical Linear Algebra with Applications, 5 (1998), pp. 33–55.
- [18] CHRIS C. PAIGE, BERESFORD N. PARLETT, AND HENK A. VAN DER VORST, *Approximate solutions and eigenvalue bounds from Krylov subspaces*, Numerical Linear Algebra with Applications, 2 (1995), pp. 115–133.
- [19] MICHAEL L. PARKS, ERIC DE STURLER, GREG MACKEY, DUANE D. JOHNSON, AND SPANDAN MAITI, *Recycling Krylov subspaces for sequences of linear systems*, SIAM Journal on Scientific Computing, 28 (2006), pp. 1651–1674.
- [20] BERESFORD N. PARLETT, *A new look at the Lanczos algorithm for solving symmetric systems of linear equations*, Linear algebra and its applications, 29 (1980), pp. 323–346.
- [21] YOUSEF SAAD, *On the Lanczos method for solving symmetric linear systems with several right-hand sides*, Mathematics of Computation, 48 (1987), pp. 651–662.
- [22] ———, *Iterative methods for sparse linear systems*, SIAM, Philadelphia, Second ed., 2003.
- [23] YOUSEF SAAD AND MARTIN H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
- [24] YOUSEF. SAAD, M. YEUNG, J. ERHEL, AND FRÉDÉRIC GUYOMARC'H, *A deflated version of the conjugate gradient algorithm*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1909–1926. Iterative methods for solving systems of algebraic equations (Copper Mountain, CO, 1998).
- [25] MARCUS SARKIS AND DANIEL B. SZYLD, *Optimal left and right additive Schwarz preconditioning for minimal residual methods with Euclidean and energy norms*, Computer Methods in Applied Mechanics and Engineering, 196 (2007), pp. 1612–1621.
- [26] VALERIA SIMONCINI, *Restarted full orthogonalization method for shifted linear systems*, BIT. Numerical Mathematics, 43 (2003), pp. 459–466.
- [27] VALERIA SIMONCINI AND DANIEL B. SZYLD, *Recent computational developments in Krylov subspace methods for linear systems*, Numerical Linear Algebra with Applications, 14 (2007), pp. 1–59.
- [28] KIRK M. SOODHALTER, DANIEL B. SZYLD, AND FEI XUE, *Krylov subspace recycling for sequences of shifted linear systems*, Applied Numerical Mathematics, In Press (2014).
- [29] ERIC DE STURLER, *Nested Krylov methods based on GCR*, Journal of Computational and Applied Mathematics, 67 (1996), pp. 15–41.
- [30] ———, *Truncation strategies for optimal Krylov subspace methods*, SIAM Journal on Numerical

- Analysis, 36 (1999), pp. 864–889.
- [31] SHUN WANG, ERIC DE STURLER, AND GLAUCIO H. PAULINO, *Large-scale topology optimization using preconditioned Krylov subspace methods with recycling*, International Journal for Numerical Methods in Engineering, 69 (2007), pp. 2441–2468.
- [32] GANG WU, YAN-CHUN WANG, AND XIAO-QING JIN, *A preconditioned and shifted GMRES algorithm for the PageRank problem with multiple damping factors*, SIAM J. Sci. Comput., 34 (2012), pp. A2558–A2575.